Chapter 1: Fundamentals of Security

Exam Objectives

- **Types of attacks**
- **Physical security**
- **Authentication and authorization**
- **Data protection**

One of the most important skills to have if you are going to support networked systems or systems connected to the Internet is the capability of securing systems and networks. And even if you are not working in a networked environment, you can apply these same skills to your customers with home Internet machines. The bottom line is that you need a solid understanding of network security.

I remember when a close friend of mine had his Web site totally replaced by a hacker. My friend's Web site files were replaced with inappropriate content, and he wondered how on Earth someone had hacked his server. It seems amazing now, but back then (circa 1994), a lot of companies did not use firewalls because they were not aware of the risks involved in having a computer connected directly to the Internet. Back then, people thought, "I have a password on the administrator account, so I am secure."

In this chapter, I introduce you to the basic concepts and terminology used to help secure an environment. Be sure to read this chapter carefully and make sure you understand the topics as you will be tested on security topics on the A+ exams. Have fun with this topic area — it is very exciting!

COPYRIGHTED MATERIAL

Identifying Types of Attacks

To me, a hacker is someone with the technical expertise to bypass the security of a network or an OS. A hacker knows how to use features of a piece of software or hardware to gain access to restricted areas of a network and then how to use those features against you and your system. For example, most Web sites connect to a database behind the scenes so that you can get a list of products when you visit their site. A hacker knows how to input

data into the site to manipulate your database server into executing the code that the hacker wants to execute — and this happens because the hacker understands the technologies being used.

The two types of hackers are

- **White-hat hackers,** who try to "hack" or break software or hardware so as to understand how to protect the environment from black-hat hackers. These are the good guys.
- **Black-hat hackers** break into a system or network for malicious reasons or for personal gain. The reasons could be for financial gain, bragging rights, or revenge.

Hackers use a number of different types of attacks to hack into a network or an OS. Sometimes an attack lays the groundwork for a future or different type of attack: that is, the initial attack does not seem all that dangerous, but it is used in the future to gain unauthorized access.

This section outlines some of the most popular types of attacks that can happen in networking environments today.

Social engineering attacks

A *social engineering attack* occurs when a hacker tries to obtain information or gain access to a system through social contact with a user. Typically, the hacker poses as someone else and tries to trick a user into divulging personal or corporate information that allows the hacker access to a system or network.

For example, a hacker calls your company's phone number, listed in the phone book, and poses as a technical support person for your company. He tells the user who answers the phone that a new application has been deployed on the network, and for the application to work, the user's password must be reset. After the password is reset to what the hacker wants, he might "verify" with the user the credential that the user uses. A user who is not educated on social engineering might divulge important information without thinking.

A social engineering attack is an attack where a hacker tries to trick a user or administrator into divulging sensitive information through social contact. After the sensitive information is obtained, the hacker can then use that information to compromise the system or network.

This example might sound unrealistic, but it happens all the time. If you work for a small company, you might not experience a social engineering attack. In a large corporate environment, though, it is extremely possible

that a social engineering attack would be successful if the company does not educate its users. A large company usually has the IT staff or management located at the head office, but most branch locations have never talked to IT management, so those branch employees would not recognize the voices of the IT folks. A hacker could impersonate someone from the head office, and the user at the branch office would never know the difference.

There are a number of popular social engineering attacks scenarios — and network administrators are just as likely to be social engineering victims as "regular" employees, so they need to be aware. Here are some popular social engineering scenarios:

- **Hacker impersonates IT administrator.** The hacker calls or e-mails an employee and pretends to be the network administrator. The hacker tricks the employee into divulging a password or even resetting the password.
- **Hacker impersonates user.** The hacker calls or e-mails the network administrator and pretends to be a user who forgot her password, asking the administrator to reset her password for her.
- **Hacker e-mails program.** The hacker typically e-mails all the users on a network, telling them about a security bug in the OS and that they need to run the `update.exe` file attached to the e-mail. In this example, the `update.exe` is the attack — it opens the computer up so that the hacker can access the computer.

Educate your users never to run a program that has been e-mailed to them. Most software vendors, such as Microsoft, state that they will never e-mail a program to a person: Instead, they will e-mail the URL to an update, but it is up to the person to go to the URL and download it. A great book to learn more on the process a hacker takes to compromise a system is Kevin Beaver's *Hacking For Dummies,* 2nd Edition (Wiley).

Network-based attacks

A *network-based attack* uses networking technologies or protocols to perform the attack. Here are the most popular types.

Ensure that you are familiar with the different types of network-based attacks for the A+ exams.

Password attacks

There are a number of different types of password attacks. For example, a hacker could perform a *dictionary attack* against the most popular user accounts found on networks. With a dictionary attack, hackers use a program that typically uses two text files:

- One text file contains the most popular user accounts found on networks, such as administrator, admin, and root.
- The second text file contains a list of all the words in the English dictionary, and then some. You can also get dictionary files for different languages.

The program then tries every user account in the user account file with every word in the dictionary file, attempting to determine the password for the user account.

To protect against a dictionary attack, be sure employees use strong passwords that mix letters and numbers. This way, their passwords are not found in the dictionary. Also, passwords are normally case sensitive, so educate users on the importance of using both lowercase and uppercase characters. That way, a hacker not only has to guess the password but also the combination of uppercase and lowercase characters.

Also remind users that words found in *any* dictionary are unsafe for passwords. This means avoiding not only English words, but also French, German, Hebrew . . . even Klingon!

Hackers can also perform a *brute force attack*. With a brute force attack, instead of trying to use words from a dictionary, the hacker uses a program that tries to figure out your password by trying different combinations of characters. Figure 1-1 shows a popular password-cracking tool known as LC4. Tools like this are great for network administrators to audit how strong their users' passwords are.

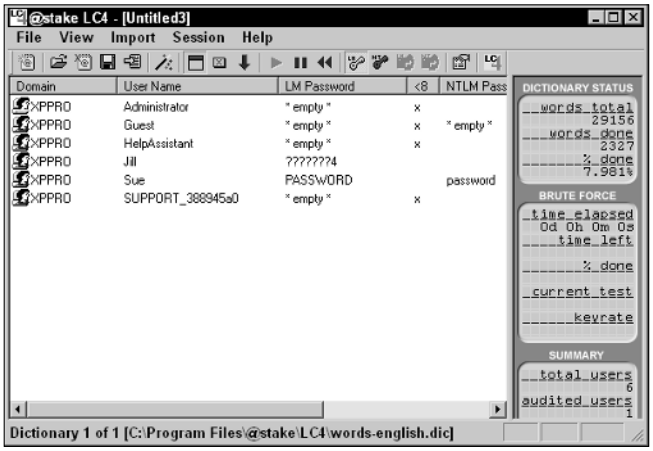

Figure 1-1: Cracking passwords with LC4.

To protect against password attacks, users should use strong passwords, which is a password comprising of letters, numbers, and symbols with a mix of uppercase and lowercase characters and a minimum length of eight characters.

Denial of service

Another popular network attack is a *denial of service (DoS)* attack, which can come in many forms and is designed to cause a system to be so busy that it cannot service a real request from a client, essentially overloading the system and shutting it down.

For example, say you have an e-mail server, and a hacker attacks the e-mail server by flooding the server with e-mail messages, causing it to be so busy that it cannot send anymore e-mails. You have been denied the service that the system was created for.

There are a number of different types of DoS attacks: for example, the ping of death. The hacker continuously pings your system, and your system is so busy sending replies that it cannot do its normal function.

To protect against denial of service attacks you should have a firewall installed and also keep your system patched.

Spoofing

Spoofing is a type of attack in which a hacker modifies the source address of a network *packet,* which is a piece of information that is sent out on the network. This packet includes the data being sent but also has a header section that contains the source address (where the data is coming from) and the destination address (where the data is headed). If the hacker wants to change "who" the packet looks like it is coming from, the hacker modifies the source address of the packet.

There are three major types of spoofing — MAC spoofing, IP spoofing, and e-mail spoofing. MAC spoofing is when the hacker alters the source MAC address of the packet, IP spoofing is when the hacker alters the source IP address in a packet, and e-mail spoofing is when the hacker alters the source e-mail address to make the e-mail look like it came from someone other than the hacker.

An example of a spoof attack is the smurf attack, which is a combination of a denial of service and spoofing. Here is how it works:

1. The hacker pings a large number of systems but modifies the source address of the packet so that the ping request looks like it is coming from a different system.

2. All systems that were pinged reply to the modified source address — an unsuspecting victim.
3. The victim's system (most likely a server) receives so many replies to the ping request that it is overwhelmed with traffic, causing it to be unable to answer any other request from the network.

To protect against spoof attacks you can implement encryption and authentication services on the network.

Eavesdropping attack

An *eavesdropping attack* occurs when a hacker uses some sort of packet sniffer program to see all the traffic on the network. Hackers use *packet sniffers* to find out login passwords or to monitor activities. Figure 1-2 shows Microsoft Network Monitor, a program that monitors network traffic by displaying the contents of the packets.

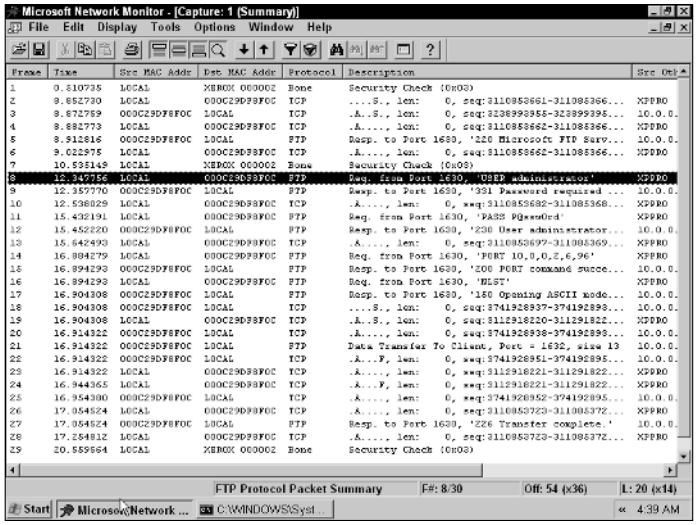
Microsoft Network Monitor - [Capture: 1 (Summary)]
File Edit Display Tools Options Window Help
Frame Time Src MAC Addr Dst MAC Addr Protocol Description Src Oth
1 0.810735 LOCAL XEROX 000002 Bone Security Check (0x03)
2 8.852730 LOCAL 080C29D78F0C TCP ....S., len: 0, seq:3110853661-311085366... XPPRO
3 8.872759 080C29D78F0C LOCAL TCP .A..S., len: 0, seq:3238998958-323899395... 10.0.0.
4 8.882773 LOCAL 080C29D78F0C TCP .A...., len: 0, seq:3110853662-311085366... XPPRO
5 8.912816 080C29D78F0C LOCAL FTP Resp. to Port 1680, '220 Microsoft FTP Serv... 10.0.0.
6 9.022975 LOCAL 080C29D78F0C TCP .A...., len: 0, seq:3110853662-311085366... XPPRO
7 10.538148 LOCAL XEROX 000002 Bone Security Check (0x03)
8 12.347756 LOCAL 080C29D78F0C FTP Req. from Port 1680, 'USER administrator' XPPRO
9 12.357770 080C29D78F0C LOCAL FTP Resp. to Port 1680, '331 Password required ... 10.0.0.
10 12.538029 LOCAL 080C29D78F0C TCP .A...., len: 0, seq:3110853682-311085368... XPPRO
11 15.432191 LOCAL 080C29D78F0C FTP Req. from Port 1680, 'PASS P@ssw0rd' XPPRO
12 15.452220 080C29D78F0C LOCAL FTP Resp. to Port 1680, '230 User administrator... 10.0.0.
13 15.642493 LOCAL 080C29D78F0C TCP .A...., len: 0, seq:3110853697-311085369... XPPRO
14 16.884279 LOCAL 080C29D78F0C FTP Req. from Port 1680, 'PORT 10,0,0,2,6,96' XPPRO
15 16.894293 080C29D78F0C LOCAL FTP Resp. to Port 1680, '200 PORT command succe... 10.0.0.
16 16.894293 LOCAL 080C29D78F0C FTP Req. from Port 1680, 'NLST' XPPRO
17 16.904308 080C29D78F0C LOCAL FTP Resp. to Port 1680, '150 Opening ASCII mode... 10.0.0.
18 16.904308 080C29D78F0C LOCAL TCP ....S., len: 0, seq:3741928937-374192893... 10.0.0.
19 16.904308 LOCAL 080C29D78F0C TCP .A..S., len: 0, seq:3112918220-311291822... XPPRO
20 16.914322 080C29D78F0C LOCAL TCP .A...., len: 0, seq:3741928938-374192893... 10.0.0.
21 16.914322 080C29D78F0C LOCAL FTP Data Transfer To Client, Port = 1632, size 13 10.0.0.
22 16.914322 080C29D78F0C LOCAL TCP .A...F, len: 0, seq:3741928951-374192895... 10.0.0.
23 16.914322 LOCAL 080C29D78F0C TCP .A...., len: 0, seq:3112918221-311291822... XPPRO
24 16.944365 LOCAL 080C29D78F0C TCP .A...F, len: 0, seq:3112918221-311291822... XPPRO
25 16.954380 080C29D78F0C LOCAL TCP .A...., len: 0, seq:3741928952-374192895... 10.0.0.
26 17.054524 LOCAL 080C29D78F0C TCP .A...., len: 0, seq:3110853723-311085372... XPPRO
27 17.054524 080C29D78F0C LOCAL FTP Resp. to Port 1680, '226 Transfer complete.' 10.0.0.
28 17.254812 LOCAL 080C29D78F0C TCP .A...., len: 0, seq:3110853723-311085372... XPPRO
29 20.559564 LOCAL XEROX 000002 Bone Security Check (0x03)
FTP Protocol Packet Summary F#: 8/30 Off: 54 (x36) L: 20 (x14)
Start Microsoft Network ... C:\WINDOWS\Syst... 4:39 AM

Figure 1-2: Using Network Monitor to analyze FTP logon traffic.

Notice in Figure 1-2 that the highlighted packet (frame 8) shows someone logging on with a username of `administrator`; in frame 11, you can see that this user has typed the password `P@ssw0rd`. In this example, the hacker now has the username and password of a network account by eavesdropping on the conversation!

To protect against eavesdrop attacks you should encrypt network traffic.

Man-in-the-middle

A *man-in-the-middle attack* involves the hacker monitoring network traffic but also intercepting the data, potentially modifying the data, and then sending out the modified result. The person the packet is destined for never knows that the data was intercepted and altered in transit.

To protect against man-in-the-middle attacks you should restrict access to the network and implement encryption and authentication services on the network.

Session hijacking

A *session hijack* is similar to a man-in-the-middle attack, but instead of the hacker intercepting the data, altering it, and sending it to whomever it was destined for, the hacker simply hijacks the conversation — a *session* — and then impersonates one of the parties. The other party has no idea that he is communicating with someone other than the original partner.

To protect against session hijacking attacks you should restrict access to the network and implement encryption and authentication services on the network.

Buffer overflow

A very popular type of attack today is a *buffer overflow attack,* which involves the hacker sending more data to a piece of software than what it is expecting. The information sent to an application is typically stored in an area of memory (a *buffer*). When more data than expected is sent to the application, the information is stored in memory beyond the allocated buffer. If the hacker can go beyond the allocated buffer, he can run their own code. This code executes in the context of the user account associated with the software that was hacked — normally an administrative account!

To protect against buffer overflow attacks you should keep the system patched and its applications.

Wireless attacks

There are a number of different attacks against wireless networks that you should be familiar with. Hackers can crack your wireless encryption if you are using a weak encryption protocol such as WEP. Hackers can also spoof the MAC address of their system and try to bypass your MAC address filters. Also, there are wireless scanners such as Kismet that can be used to discover wireless networks even though SSID broadcasting is disabled.

To protect against wireless attacks you should implement encryption protocols such as WPA2 and use an authentication server such as a RADIUS server for network access For more information on wireless check out Book VIII, Chapter 2.

Software-based attacks

Just as there are a number of different types of network attacks, there are a number of software attacks as well. As you can likely guess, a *software attack* comes through software that a user runs. The most popular software attacks are mentioned in the sections that follow, and you should be familiar with them for the A+ exams.

SQL injection

A *SQL injection attack* occurs when the hacker sends `Transact SQL` statements (statements that manipulate a database) into an application so that the application will send those statements to the database to be executed. If the application developer does not validate data inputted in the application, the hacker can modify the data or even delete it. The hacker can potentially manipulate the OS through the application that sends the input to the database.

Trojan horse

A *Trojan horse* is software that a user is typically tricked into running on the system; and when the software runs, it does something totally different than what the user expected it to do. For example, NetBus (an older attack) is an example of a Trojan horse program sent as a file called `patch.exe`. The user receiving the file, typically through an e-mail, believes that the file will fix a security issue. The problem is that `patch.exe` is a Trojan horse, and when that horse starts running, it opens the computer up to allow a hacker to connect to the system.

The hacker then uses a client program, like the one shown in Figure 1-3, to connect to the system and start messing with the computer. The hacker can do things like launch other programs, flip your screen upside-down, eject your CD-ROM tray, watch your activity, and modify or delete files!

Virus

A *virus* is a program that causes harm to your system. Typically, viruses are spread through e-mails and are included in attachments, such as word processing documents and spreadsheets. The virus can do any of a number of things: delete files from your system, modify the system configuration, or e-mail all your contacts in your e-mail software. To prevent viruses, install antivirus software and do not open any file attachments that arrive in your e-mail that you are not expecting.

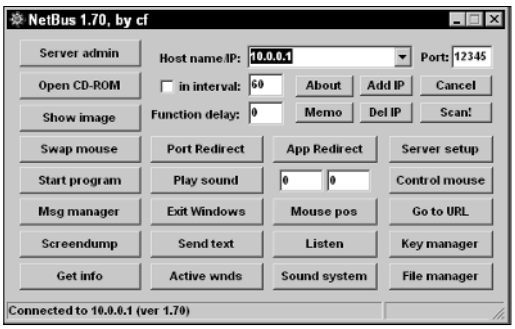

Figure 1-3: Using NetBus to control a user's computer.

Worm

A *worm* is a virus that does not need to be activated by someone opening the file. The worm is *self-replicating,* meaning that it spreads itself from system to system, infecting each computer. To protect against a worm, you should install a firewall. A *firewall* is a piece of software or hardware that prevents someone from entering your system.

Logic bomb

A *logic bomb* is malicious software that could run every day, but the software was designed to wreak havoc on your system on a certain date and time. The scary thing about logic bombs is that they seem like useful software until the day the programmer decides it will become malicious!

To protect against malicious software such as a virus, Trojan horse, worm, and a logic bomb you should use a firewall and keep your virus definitions up-to-date.

Understanding Physical Security

You should implement security in many places, and one of the most overlooked areas is physical security. *Physical security* has nothing to do with software; rather, it covers how you protect your environment and systems by making sure that a person cannot physically access the system. For example, many companies use a numeric keypad to secure entrance to a facility. To get into the facility, users must enter a valid combination to open the door.

Another example of physical security is the server room. Most server room doors are locked with a numeric padlock or a key. Higher-security server rooms sometimes even require fingerprint or retinal scans from anyone trying to enter the room. The benefit of locking your servers in the server

room is a hacker cannot boot off a bootable CD-ROM, which could bypass the OS entirely. After a hacker bypasses the OS, he typically can bypass a lot of the security by booting to a totally different OS.

You can apply enterprise security best practices to your home systems. For example, to help secure your home system, you might want to prevent booting from a CD-ROM so that an unauthorized person cannot try to bypass your Windows security.

BIOS settings

You can set a number of settings in your system BIOS to help control the security of the system. Be sure to investigate the BIOS settings on your system to see what security settings you can enable on the system. Here are some popular BIOS/CMOS settings to aid in physical security:

- **Drive lock:** Drive lock (a popular feature with laptops) is a hard disk specification used to protect access to the drive. To protect access to the drive, there are two drive lock passwords: a user password and a master password. The user password is used by the user wanting to access the system; the master password is used to reset the user password if the user forgets the password. Do not confuse drive lock passwords that prevent booting from the drive with the general CMOS passwords for the system. If the user password and master password are forgotten or lost, the drive is useless.
- **Passwords:** You can set a power-on password in CMOS to limit who can use the system. If the power-on password is forgotten, it can typically be erased via a jumper on the motherboard or by taking the battery off the motherboard and putting it back in.
- **Intrusion detection:** Most systems have intrusion detection features that can be enabled through the BIOS that will notify you if the cover is taken off the system. This is designed to alert you if someone opens the cover and takes internal components.
- **TPM:** The *Trusted Platform Module* (TPM) is a chip on computer hardware used to store cryptography keys that are typically used to encrypt data. A TPM chip can also be used to authenticate a device because it contains a unique key that identifies the chip, or hardware device. Most computers today have a TPM chip, and a number of software solutions (such as Windows Vista Bit Locker) can use the TPM chip to encrypt the contents of the drive.

Best practices

To protect your systems, follow these physical security best practices:

- **Secure server placement.** Lock your servers in a room for which only a select few individuals have the key.
- **Lock the workstation.** When you leave your system, get in the habit of locking your workstation. A locked workstation can only be unlocked by yourself or the network administrator. This will prevent other users from accessing the system while you are away.
- **Disable boot devices.** Dsable the ability to boot from a floppy disk or CD-ROM in the CMOS setup on the systems.
- **Set CMOS password.** Because most hackers know how to go to CMOS and enable booting from CD-ROM, make sure that you set a password on CMOS so that a hacker cannot modify your CMOS settings. Figure 1-4 shows a CMOS password being enabled.

  Check out Book II, Chapter 4, to get the lowdown on reconfiguring your CMOS settings.

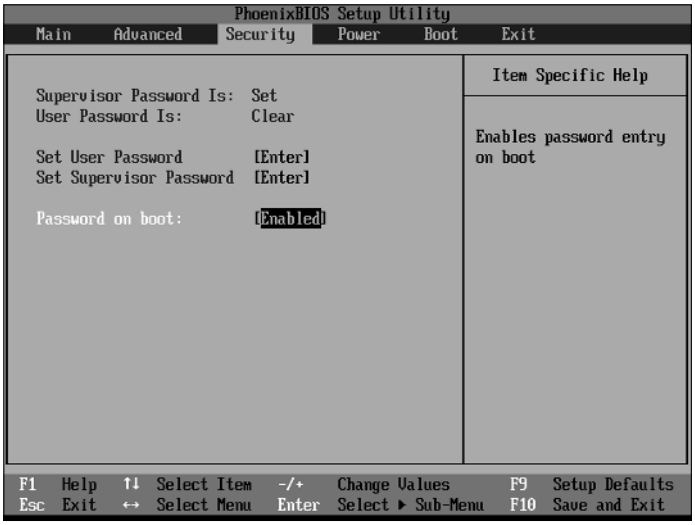
PhoenixBIOS Setup Utility
Main Advanced Security Power Boot Exit
Item Specific Help
Supervisor Password Is: Set
User Password Is: Clear
Enables password entry on boot
Set User Password [Enter]
Set Supervisor Password [Enter]
Password on boot: [Enabled]
F1 Help ↑↓ Select Item -/+ Change Values F9 Setup Defaults
Esc Exit ↔ Select Menu Enter Select ▸ Sub-Menu F10 Save and Exit

Figure 1-4: Enabling the CMOS password.

- **Disable network ports.** To prevent a hacker from entering your office, plugging into the network, and performing a number of network attacks, ensure that network ports in lobbies and front entrances are disabled unless an administrator enables them.
- **Use a lockdown cable.** Use a *lockdown cable* to tether laptops, projectors, and other types of office equipment to a table or desk. Figure 1-5 shows a lockdown cable being used to secure a laptop. A lockdown cable usually connects to a special hole in the side of the computer equipment (look for a picture of a lock next to it).

Figure 1-5: A lockdown cable is used to secure computer equipment to a desk.

Remembering ways to physically secure your systems will help you with the security portion of the A+ exam. Be sure to place critical systems in locked rooms and lock down equipment that is accessible by the public.

Understanding Authentication and Authorization

After you physically secure your environment, focus on the people who access your systems and network. The next step after implementing physical security is to ensure that persons who enter your server room or have a connection to a network port are authorized to log on to the network. Logging onto the network is *authentication.*

Authentication

Authentication is the process of proving one's identity to the network environment. Typically, authentication involves typing a username and password on a system before you are granted access, but you could also use biometrics to be authenticated. *Biometrics* is using one's unique physical characteristics, such as a fingerprint or the blood vessels in one's retina, to prove one's identity. Figure 1-6 shows a fingerprint reader used to scan your fingerprint when logging on.

Here is a quick look at what happens when you log on to your system with a username and password. When you type a username and password to log on to a system, that username and password are verified against a database —

the *user account database* — which has a list of the usernames and passwords allowed to access the system. If the username and password you type are in the user account database, you are allowed to access the system. Otherwise, you get an error message and are not allowed access.

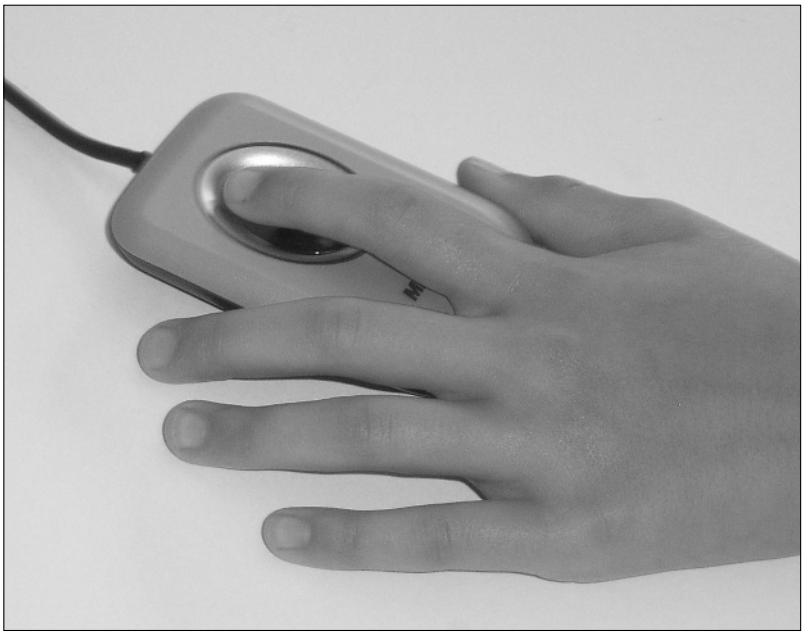

Figure 1-6: A fingerprint reader is an example of biometrics used for authentication.

The name of the account database that stores the usernames and passwords is different, depending on the environment. In a Microsoft network, the account database is the *Active Directory Database* and resides on a server known as a *domain controller* (shown in Figure 1-7).

Generating the access token

When you log on to a Microsoft network environment, the username and password you type are placed in a logon request message that is sent to the domain controller to be verified against the Active Directory Database. If the username and password that you typed are correct, an access token is generated for you. An *access token* is a piece of information that identifies you and is associated with everything you do on the computer and network. The access token contains your user account information and any groups of which you are a member. When you try to access a resource on the network, the user account and group membership in the access token are compared against the permission list of a resource. If the user account in

the access token or one of the groups contained in the access token are also contained in the permission list, you are granted access to the resource. If not, you get an access-denied message.

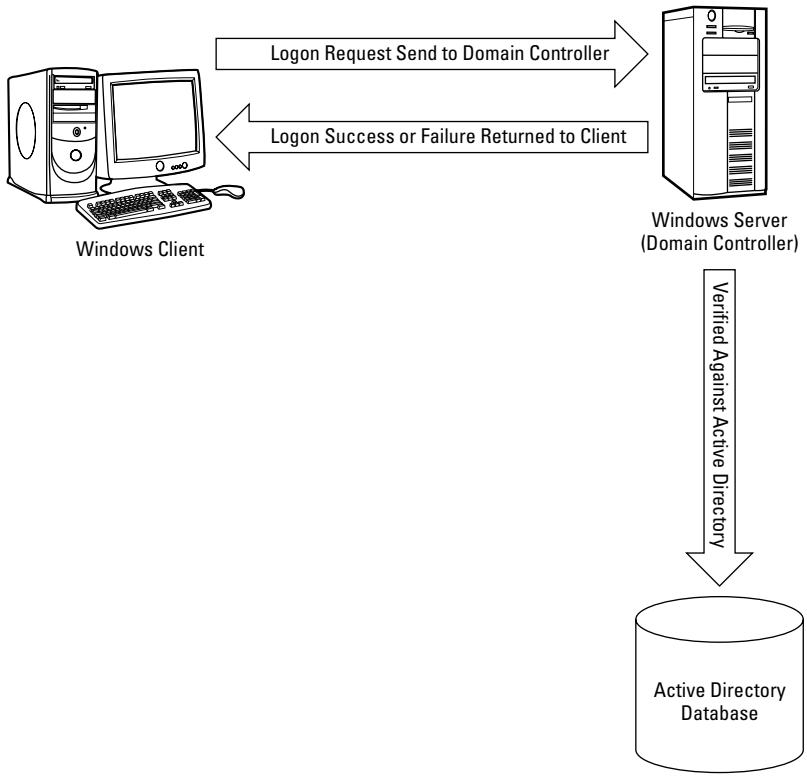

Figure 1-7: Logging on to Active Directory in a Microsoft network environment.

If you do not have a server-based network environment and you are simply running Windows 2000 Professional or Windows XP, when you log on, the logon request is sent to the local computer — to an account database known as the Security Accounts Manager (SAM) database. When you log on to the SAM database, an access token is generated as well, and that helps the system determine what files you can access.

Smart card

Another type of logon supported by network environments today is the use of a smart card. A *smart card* is a small, ATM card–like device that contains your account information. You insert the smart card into a smart

card reader that is connected to a computer, and then you enter the PIN (personal identification number) associated with the smart card. This is an example of securing an environment by forcing someone to not only have the card but also know the PIN.

Strong passwords

It is really hard to talk about authentication without talking about ensuring that users create strong passwords. A *strong password* is a password that is very difficult for hackers to guess or crack because it contains a mix of uppercase and lowercase characters, contains a mix of numbers and letters, and is a minimum of eight characters long.

Authorization

After a user logs on and an access token is created, the user may start trying to access resources such as files and printers. To access a file, folder, or printer on the network, the user must be authorized to access the resource. *Authorization* is the process of giving a user permission to access a resource or the right to perform an OS task. Do not confuse authentication and authorization: You must be first authenticated to the network; then, after authentication, you can access the resources you have been authorized for.

Permissions

To authorize access to a resource, you set permissions on the resource. For example, if you want to allow Jill to access the accounting folder, you need to give Jill permission to the accounting folder, as shown in Figure 1-8.

Using strong passwords

A number of years ago, I had a co-worker who was always trying to get me to guess his passwords. He thought I had some magical trick or program that was cracking them, but all I was doing was guessing his passwords. I remember one time he changed it, and I could not guess it until one night when we were at a social function for work and all he talked about were the Flyers hockey team. I remember sitting there thinking, "I bet that is his password." Sure enough, the next day at work, I tried `flyers` as his password, and it worked. Now the lesson here is that he should have at least mixed the case of the word *flyers* to make something like `flYeRs`, or even better, thrown a symbol in there by replacing the "s" with a "$." I would have had a much harder time trying to guess his password if he had used `flYeR$` instead. This is an example of a strong password.

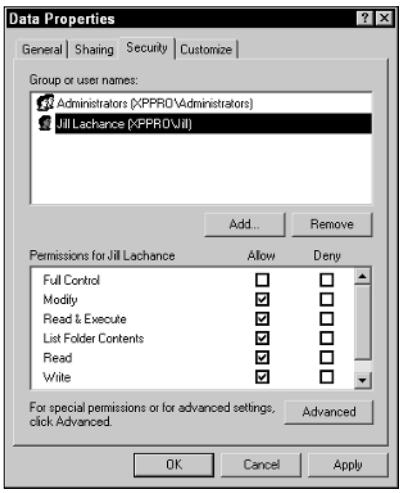

Figure 1-8: Using permissions to authorize which users are allowed to access the resource.

In Figure 1-8, you can see that the Administrators and Jill have access to the resource. No one else is authorized to access the resource. You find out how to set permissions in the next chapter, but for now, make sure you understand the difference between authentication and authorization.

Rights

In the Windows world, there is a difference between permissions and rights. As you can read in earlier sections, permission is your level of access to a resource. Comparatively, a right is your privilege to perform an OS task. For example, you can be assigned the right to change the time on the computer. Other examples of rights are the right to do backups or the right to log on to the system.

To learn more about how to set permissions and rights, check out Book IX, Chapter 2.

Methods of Securing Transmissions

After you authenticate users and authorize them to access certain parts of the network, you should then consider methods of securing information while it travels along the network cable.

Most network communication is sent along the network wire in *cleartext,* meaning that anyone connected to your network can read the information. But if the information is traveling across the Internet, anyone can view that information if it is passed in cleartext.

Most Internet protocols, such as HTTP, send information in cleartext, and it is up to the people who set up the servers that use these Internet protocols to encrypt the information before it is released to the Internet. *Encrypting* the information means that the information is run through a mathematical calculation that generates an altered version of the information: a *result.* For example, the words `Glen Clarke` could be encrypted to look like `7y3i s3fk4r`. If anyone intercepts such encrypted information and views it while it is traveling across the wire, the information would mean nothing.

Here is a real-world example. You type your credit card number on a Web site, but you certainly do not want that credit card number to be viewed while you send it from your client computer to the server. You want to be sure that the Web site where you enter the credit card number encrypts the traffic. You can tell by the lock icon that appears in the Web browser, as shown in Figure 1-9.

login - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Microsoft Learning
MCP Member Home | MCT Member Home | Learning Home
Welcome to the Microsoft Certification Member Sites.
You must be a certified professional and have a .NET Passport to proceed.
Please click the Passport Icon Sign In to sign in or create a Passport.
Internet
Start
6:02 AM

Figure 1-9: Identifying a secure site by locating the lock in Internet Explorer.

It is important for the A+ exam that you understand popular methods of encrypting traffic. You can use a number of technologies, such as

- **Secure Sockets Layer (SSL):** This protocol is used to encrypt different types of Internet traffic. For example, you could use SSL to encrypt HTTP traffic by applying a digital certificate to the Web site. The *digital certificate* contains the key that is used to encrypt and decrypt the traffic.
- **Internet Protocol Security (IPSec):** This protocol can encrypt all TCP/IP traffic between systems. As a network administrator, you configure IPSec on the server and the clients with the same key, which is used to

encrypt and decrypt network traffic. Because of the configuration, it is an unlikely solution for a Web site but is a great way to encrypt traffic on your network.

- **Virtual Private Network (VPN):** A VPN allows a user to connect across the Internet to a remote network, typically her office network, and send information between her system and the office network securely. The information is secured because the VPN technology used creates an encrypted tunnel between the user and the office network — any data that travels through the tunnel is encrypted.

The preceding sections touch on a number of places that require security. Here is a quick overview of the security steps that I have discussed so far:

- Secure your office environment first from physical access by unauthorized persons.
- Set up a system for authentication, which is the idea that users must log on to the network.
- After users log on to the network, they must be authorized to access resources.
- When you allow someone to access resources, make sure that you encrypt the traffic while it is in transit, especially if the information is transmitted outside your own network.

Do Not Forget about Data Protection

In this section, you find out about how to secure your data environment from a hacker or malicious user. When securing your systems, you want to protect the systems from a person who damages information or systems with or without intent. You want to be sure to secure your environment from hackers, but at the same time, you want to protect your systems from users on the network who may cause damage without meaning to. Accidents can happen, so be sure to prevent accidents from happening by following the best practices in the following sections.

Destroying data

Most office environments have strict policies in place to help secure confidential information. Shredding paper documents with personal or confidential information is a no-brainer, and computerized data should be no different. A company needs strict guidelines on how to destroy data that resides on computer hard drives.

Hard drive destruction

Destroying data that resides on a computer hard disk typically involves shredding the computer hard drive with a huge shredding machine, or destroying the drive another way, such as sanding the platters down to nothing.

I have talked to some customers who used to destroy drives by driving spikes through them, but what they found was that the data around the hole that the spike put in the drive could still be read! These customers now disintegrate the drive in a huge shredder. Other customers sand the drives right down to nothing. Either way, if securing the data is a concern, make sure to physically destroy the disk that contains the data.

Data wiping

Instead of destroying the drives with a traditional shredded, some companies use a *shredder application,* which writes a bunch of 1s to the drive, thereby overwriting the previous data. These applications typically overwrite the drive a number of times because hackers can retrieve the data from disk even after it has been overwritten a few times. When purchasing shredding software, be sure to investigate how many overwrite operations the software performs. I recommend using software that overwrites at least seven times.

Hard drive recycling

What do you do with your hard drive when you get an upgrade or replacement drive? Well, you could pass the hard drive on to someone else who needs the drive, or you could have the drive recycled. After all, a number of useful parts are on the drive, such as the hard drive platters and the magnets.

However, be aware of the risk of passing your drive over to someone! Companies concerned with corporate security and data privacy will likely opt to destroy the drive instead of recycling because of the risk of having private data lifted off the drive.

Backing up data

A big part of securing the data environment is not only setting the permissions but also ensuring that you create a good backup and restore strategy. Identify which files are critical to the operation of the business and should be backed up. You also want to be familiar with all types of information used by your company. For example, you might depend on e-mail, so make sure that you back up your e-mail server along with any files in shared folders. If your company stores important data in databases, make sure that you back up those databases as well.

Backup review

You can find out more about backups in Book VII, Chapter 3, but for the exam, here are some of the key points you need to remember.

When you perform a backup, the OS keeps track of which files have been changed since the last backup by setting the *archive bit.* The archive bit is an attribute of the file that tells the system that the file has changed. To view the archive bit within Windows XP or Vista, right-click the file and choose Properties. In the Properties dialog box, click the Advanced button.

The first option in the Advanced Attributes dialog box that appears — File Is Ready for Archiving — is the archive bit. (See Figure 1-10.) When this check box is selected, it means that the file needs to be backed up because it has changed.

Advanced Attributes
Choose the options you want for this file.
Archive and Index attributes
File is ready for archiving
For fast searching, allow Indexing Service to index this file
Compress or Encrypt attributes
Compress contents to save disk space
Encrypt contents to secure data
Details
OK
Cancel

Figure 1-10: Viewing the archive bit in Windows.

Before you perform a backup, decide what type of backup to perform. Each backup type deals with the archive bit a little differently. Here are the three major types of backup:

- **Full backup:** Copies any files that you select, whether the archive bit is set or not, and clears the archive bit on any file that is backed up — essentially recording the fact that the file has been backed up.
- **Differential backup:** Copies any files that have changed, but it does not clear the archive bit; thus, there is no record that the files have been backed up. The benefit is that the next time you do the backup, the files will be backed up again because the archive bit has not been cleared. As far as the OS is concerned, the file has not been backed up since it was changed.
- **Incremental:** Copies any file that has changed and then clears the archive bit on any files that are backed up. Thus, if a file is copied

during an incremental backup, because the backup process clears the archive bit, the file will not be backed up during subsequent incremental backups unless the file changes again.

Be familiar with the difference between a full, incremental, and differential backup. Also know which backup types clear the archive bit.

Tape rotation and offsite storage

Take the time to rotate tapes so that the same tape is not used all the time. You also want to make sure that you store a backup offsite in case of a disaster such as flood or fire. It is important that you are able to recover the system no matter what happens.

Test restore operations

As a last point with backup strategy best practices, you want to test restorations frequently to ensure that you can recover information from backup without any problem. You do not want to find out that the backups are bad when management is hanging over your shoulder waiting for the company network to come back online! Be sure to perform regular test restorations.

Implementing RAID solutions

To help secure your data, not only do you want to have good backups, but you also want to ensure that you are implementing some form of a RAID solution. RAID (Redundant Array of Inexpensive Disks) is covered in detail in Book II, Chapter 5, so in this section I review the different types of RAID volumes supported in Windows servers and ensure that you understand that RAID solutions are a way of helping secure data.

RAID is a way of storing duplicated data on multiple disks; if one disk goes down, the data is still available to the users because other disks in the RAID array have a copy of the data. The benefit of using RAID instead of backups is that with the RAID solution, the user never knows that a drive has failed because the other drive is supplying all the data. ***Note:*** You still need the backups, though, in case both drives fail, or some disaster happens, like a flood or fire, destroying the system and all of its drives.

A number of different types of RAID solutions are available. The ones provided by the Windows Server OSes are as follows:

- **RAID Level 0:** Also known as a *striped volume* in Windows, RAID Level 0 writes different parts of the data to different disks at the same time. The benefit of a striped volume is that you get a performance benefit

by writing the data at the same time to two different disks, essentially taking less time to read or write to the file. Note that the data is split between both drives, and there is no duplication — which means that this is not really a redundant solution.

- **RAID Level 1:** Also known as *a mirrored volume* in Windows. A mirrored volume duplicates the data stored on one disk to another disk. If one disk fails, the other disk has a copy of the data.
- **RAID Level 5:** Also known as a *RAID 5 volume* in Windows. A RAID 5 volume requires a minimum of three drives and writes to all drives in the solution like a striped volume. A RAID 5 volume is different than a striped volume in the sense that it does store redundant data — *parity data* — on one of the disks. The redundant data is used to calculate the missing data when a disk goes missing, ensuring that users can still retrieve the data without noticing a problem.

Be comfortable with the RAID levels when preparing for the exam. Check out Book II, Chapter 5, to see how to create volumes in Windows.

Data encryption

Encrypting data converts information to an unreadable format so that if folks gain access to the data, they cannot understand it. In the cryptography world, encryption is described as changing plain text to cipher text. As you can likely intuit, decryption converts cipher text to plain text.

There are a number of ways to encrypt data on the hard drive:

- **EFS:** The Encrypting File System (EFS) is a feature of NTFS and can be enabled through the file properties. After the file is encrypted, it can be read only by authorized persons. To read more about EFS, check out Book V, Chapter 4.
- **Bit Locker:** Instead of encrypting data at the file level with EFS, you can have Windows encrypt the entire partition or volume, which protects all data on the partition, including the Windows OS, the Registry, and the data. With Bit Locker, data is encrypted by using keys stored in a TPM chip or a USB drive, depending upon how Bit Locker has been configured.
- **Third-party software:** You can also use third-party software to encrypt data. For example, you can use the free program TrueCrypt (`www.truecrypt.org`) to encrypt all your data into a TrueCrypt file and then copy the file to a USB flash drive.

Compliance and classification

Part of securing your computing environment is understanding how to protect the business from disclosure of information and by ensuring the business is compliant with government and industry regulations that surround computing and data.

Compliance

Data compliance is the concept of protecting the data from information leaks and ensuring recoverability of information by following government regulations and industry regulations. For example, if you are in the health industry, you must be compliant with the Health Insurance Portability and Accountability Act (HIPAA), which requires that health records and patient information be secured and kept private.

Companies are also required to protect customer information as outlined by to the Privacy Act. Companies are no longer allowed to share customer information with other businesses, including contact information such as e-mail addresses and phone numbers.

Another example of data compliance is ensuring that your company has taken the correct steps and implemented the correct controls to adhere to the Sarbanes-Oxley Act, which outlines that the company must be able to prove that adequate auditing controls have been put in place in case an incident requires review of internal information, such as company e-mails.

If you are the security manager for your company, be sure to spend time researching which government regulations and industry regulations your company falls under. With your list of regulations in hand, then you can determine the steps you need to take to be compliant.

Classification

Part of securing company information is through data classification, which assigns a level of sensitivity to information, such as Confidential or Top Secret. After the level of sensitivity is assigned to the information, the necessary controls are put in place to protect that classification of information.

Each data classification has specific security measures that need to be implemented to keep it secure. For example, a company might decide that top secret information cannot leave the "top secret" system — say, by disabling the ports on the system that typically would allow connecting a removable drive.

Data classification is assigned to the information based on the value of the information to the organization. Each classification level is designed to indicate whether the information is to be kept private or is available for public release. The following are examples of classification levels:

- Top secret, secret, and unclassified
- Confidential, official use only, and public
- Highly confidential, proprietary, internal use, and public

Getting an A+

This chapter introduces you to a number of security-related terms that you need to understand before taking your first A+ exam. Here some key points to remember when preparing for the exam:

- *Authentication* is the process of proving an identity to the network, but *authorization* is the process of determining whether accessing a resource is allowed after authentication takes place.
- Hackers take many different approaches to compromise a system. Protect your environment from both network-based and software-based attacks, and make sure that physical security is in place.
- A *denial of service* (DoS) is an attack on a system or network that prevents the system or network from performing its regular function.
- *Social engineering* is a popular type of attack that involves the hacker compromising security by tricking an employee through social contact. The social engineer might entice the user to divulge confidential information or might trick the user into running a program that does harm to the system.
- You secure network traffic by *encrypting* traffic between two systems by using technologies such as SSL and IPSec. Administrators typically use SSL to encrypt Web traffic and IPSec to encrypt internal or VPN traffic.
- Securing your data involves not only protecting resources with permissions but also protecting your data by following proper data destruction procedures and backup strategies as well as creating redundant disk solutions.

Prep Test

1 **What type of attack involves the hacker tricking a user through social contact?**

A ❍ Password attack
B ❍ Eavesdrop attack
C ❍ Man-in-the-middle attack
D ❍ Social engineering attack

2 **What type of attack involves the hacker using a packet sniffer and trying to view confidential information traveling over the network?**

A ❍ Password attack
B ❍ Eavesdrop attack
C ❍ Man-in-the-middle attack
D ❍ Social engineering attack

3 **What type of attack involves the hacker causing your system or network to become unresponsive to valid requests?**

A ❍ DoS attack
B ❍ Eavesdrop attack
C ❍ Man-in-the-middle attack
D ❍ Password attack

4 **What type of RAID volume duplicates the data fully on two disks?**

A ❍ Striped volume
B ❍ Mirrored volume
C ❍ RAID 5 volume
D ❍ RAID Level 0

5 **What should you purchase for each laptop to help protect it from theft?**

A ❍ Flash drive
B ❍ Driver disk
C ❍ A Doberman Pinscher
D ❍ Lockdown cable

6 **Which of the following are forms of biometrics? (Select all that apply.)**

A ❑ Fingerprint scan
B ❑ Smart card
C ❑ Username and password
D ❑ Retinal scan

7 **What type of backup copies the files that have changed but does not clear the archive bit?**

A ❍ Full backup
B ❍ Incremental backup
C ❍ Differential backup
D ❍ Copy

8 **What technology is typically used to encrypt traffic between a Web server and Web browser?**

A ❍ DoS
B ❍ IPSec
C ❍ Smart card
D ❍ SSL

9 **In high-security environments, what should you do with old hard drives?**

A ❍ Donate them to charity.
B ❍ Recycle them.
C ❍ Physically destroy them.
D ❍ Drive a spike through them.

10 **What type of attack involves the hacker modifying the source address of the packet?**

A ❍ Spoof attack
B ❍ Eavesdrop attack
C ❍ Man-in-the-middle attack
D ❍ Social engineering attack

Answers

1 **D.** Social engineering is a type of hack that involves contacting victims through phone or e-mail and tricking them into doing something that compromises company security. *See "Social engineering attacks."*

2 **B.** An eavesdropping attack occurs when a hacker monitors network traffic to try to capture information that could be useful in another attack. *Review "Eavesdropping attack."*

3 **A.** A denial of service (DoS) attack is when a hacker consumes all the system's processing power or bandwidth so that it cannot to perform its normal job. *Check out "Denial of service."*

4 **B.** A mirrored volume is used to create a full duplicate of the data on two different disks. *Peruse "Implementing RAID solutions."*

5 **D.** A lockdown cable is used to secure the laptop to a desk to help prevent the laptop from being stolen. *Refer to "Understanding Physical Security."*

6 **A, D.** Biometric devices involve authenticating a user through the user's unique physical characteristics. Fingerprint scans and retinal scans are popular biometric authentication methods. *See "Authentication."*

7 **C.** A differential backup backs up only those files that have changed since the last full backup and then does not clear the archive bit. *Check out "Backup review."*

8 **D.** Secure Socket Layer (SSL) is used to encrypt Web traffic. You can identify whether you are on a secure Web site by looking for the lock icon at the bottom of the screen. *Peruse "Methods of Securing Transmissions."*

9 **C.** You want to make sure that you physically destroy the drives if securing data is critical to the business. *Take a look at "Destroying data."*

10 **A.** A spoof attack occurs when the hacker modifies the source address, trying to hide the origin of the packet. *Refer to "Spoofing."*